Geometric scaling: a simple preconditioner for certain linear systems with discontinuous coefficients

Dan Gordon* Rachel Gordon†

April 17, 2009

Abstract

Linear systems with large differences between coefficients ("discontinuous coefficients") arise in many cases in which partial differential equations (PDEs) model physical phenomena involving heterogeneous media. The standard approach to solving such problems is to use domain decomposition (DD) techniques, with domain boundaries conforming to the boundaries between the different media. This approach can be difficult to implement when the geometry of the domain boundaries is complicated or the grid is unstructured. This work examines the simple preconditioning technique of scaling the equations by dividing each equation by the L_p-norm of its coefficients. This preconditioning is called geometric scaling (GS). It has long been known that diagonal scaling can be useful in improving convergence, but there is no study on the general usefulness of this approach for discontinuous coefficients. GS was tested on several nonsymmetric linear systems with discontinuous coefficients derived from convection-diffusion elliptic PDEs with small to moderate convection terms. It is shown that GS improved the convergence properties of restarted GMRES and Bi-CGSTAB, with and without the ILUT preconditioner. GS was also shown to improve the distribution of the eigenvalues by reducing their concentration around the origin very significantly.

Keywords. *Bi-CGSTAB; diagonal scaling; discontinuous coefficients; domain decomposition; geometric scaling; GMRES; GS; Lp-norm; linear equations; nonsymmetric equations; parallel processing; partial differential equations.*

arXiv:0812.2769v2 [cs.MS] 4 May 2009

1 Introduction

Many physical phenomena are modeled by partial differential equations (PDEs) which describe the relations between one or more scalar or vector fields and the surrounding media. When boundary conditions are prescribed, a common approach to achieving a numerical solution is to impose a grid and discretize the equations, thus getting a system of linear equations. In some cases, this approach yields a system of equations with very large differences between coefficients of the equations. Examples of such systems arise in modeling flow through heterogeneous media with widely-varying permeability, oil reservoir simulation, electromagnetics and semiconductor modeling. Such systems are commonly referred to as systems with "discontinuous coefficients".

*Corresponding author. Dept. of Computer Science, University of Haifa, Haifa 31905, Israel. gordon@cs.haifa.ac.il

†Dept. of Aerospace Engineering, The Technion–Israel Inst. of Technology, Haifa 32000, Israel. rgordon@tx.technion.ac.il

One of the most common methods for tackling such problems is the domain-decomposition (DD) approach, in which the domain is partitioned into subdomains, with the subdomain boundaries conforming to the boundaries between the different media. DD techniques typically operate as follows: some boundary conditions are assumed to exist on the interfaces between subdomains, and a solution of the equations in each subdomain is obtained, often with only low accuracy. The boundary conditions at the interfaces are then updated according to these solutions, and the process is repeated until convergence is achieved. There exists a vast literature on this subject, and a review of the area is beyond the scope of this work; see, for example, [4, 18, 13, 14]. DD techniques may be difficult to implement on unstructured grids or when the boundaries between domains have a complicated geometry.

We consider nonsymmetric linear systems with discontinuous coefficients derived from convection-diffusion elliptic PDEs with small to moderate convection terms. It is shown that a simple preconditioning technique, which we call *geometric scaling* (GS), can be applied to the system of equations in order to improve the convergence properties of algorithms applied to the system. GS(p), with an integer parameter $p \geq 1$, consists of dividing each equation by the L_p-norm of its vector of coefficients. GS, which is a particular form of diagonal scaling on the left, is geometric in the following sense: after applying GS(p) to a linear system, all the rows of the system matrix have a unity L_p-norm.

In order to examine the general usefulness of GS in conjunction with various solution methods, we tested it with the two leading Krylov subspace solvers for nonsymmetric systems: GMRES [17] and Bi-CGSTAB [21]. Both algorithms were tested with and without the ILUT preconditioner [15] on several nonsymmetric problems taken from (or based on) Saad [16], van der Vorst [21], and Graham and Hagger [8]. Four basic problems were considered, as well as several variations of the problems, such as modifications of the differences between the coefficients and various grid sizes. Results are provided for three different levels of accuracy.

Our experiments used the L_2-norm, but the L_1-norm yielded similar results. Our results can be characterized as follows:

- In most cases, when the tested method (algorithm/preconditioner combination) converges to the specified accuracy criterion, GS speeds up the convergence time.
- In many cases, when the tested method stagnates or diverges on the original system, it converges on the scaled system.
- GS was not particularly useful on problems without discontinuous coefficients.
- GS was not helpful on problems derived from PDEs with large convection terms. Such PDEs produce linear systems with very large off-diagonal elements, and these require other solution methods.

We also provide data on the effect of GS on the distribution of the eigenvalues. It is generally considered to be detrimental to the convergence of Krylov subspace methods to have a large accumulation of eigenvalues near the origin. In the cases examined here, there is a large accumulation of eigenvalues near the origin, and GS appears to "push" many eigenvalues away from the origin. Also, the eigenvalues of the scaled system appear to be distributed similarly to those obtained without scaling but with equal-sized coefficients.

Note that the above choice of algorithm/preconditioner combinations does not imply that these methods are optimal for the above problems. Finding an optimal method is problem-dependent and a topic for further research.

Diagonal scaling is not new, and its usefulness for discontinuous coefficients has also been noted. Van der Sluis [20] deals with the effect of scaling on the condition number of matrices. Widlund [22, p. 34-35] writes that well-scaled ADI methods give good rates of convergence when the coefficients of elliptic problems vary very much in magnitude (ADI stands for the alternating direction implicit method [12]). Graham and Hagger [8, p. 2042-2043] write that diagonal scaling has been observed in practical computations to be very effective as a preconditioner for problems with discontinuous coefficients. Duff and van der Vorst [2] write that on vector machines, diagonal scaling is often competitive with other approaches. Regarding diagonal scaling of symmetric matrices, see Meurant [11, Th. 8.1, 8.2], who also has some reservations about the usefulness of diagonal scaling for parallel processing [11, p. 401]. Gambolati et al. [3] use the least square logarithm (LSL) scaling on the rows and the columns of the system matrix for a certain problem in geomechanics with discontinuous coefficients.

However, most previous works are concerned with symmetric systems and diagonal scaling refers to multiplying the system matrix on the left and on the right by the same diagonal matrix, in order to preserve symmetry. Also, there is no study on the general usefulness of geometric scaling for nonsymmetric problems with discontinuous coefficients. The idea of geometric scaling was found to be useful for certain problems in image reconstruction from projections; see [7].

DD methods are often mentioned in connection with parallel processing. The equations of different domains can, in principle, be solved in parallel by different processors. However, the different domains that arise in many practical situations do not necessarily lead to an optimal load-balancing assignment of equations to processors. With the GS approach, inherently parallel algorithms such as Bi-CGSTAB and GMRES can be implemented efficiently in parallel without regard to subdomain boundaries. Needless to say, GS is inherently parallel.

The rest of this paper is organized as follows. §2 presents some essential background. §3–§6 deal with the four different problems, and §7 concludes with a summary and some future research directions.

2 General background

Throughout the rest of the paper, we assume that all vectors are column vectors, and we use the following notation: $\langle p, q\rangle$ denotes the dot product of two vectors p and q, which is also $p^T q$. Given a vector $x = (x_1, \ldots, x_n)^T \in \mathbb{R}^n$, we denote its L_p-norm by $\|x\|_p = (x_1^p + \cdots + x_n^p)^{1/p}$. For $p = 2$, we will omit the index and just write $\|x\| = \|x\|_2 = \sqrt{\langle x, x\rangle}$. If A is an $n \times n$ matrix, we denote by a_i the ith row-vector of A; i.e., $a_i = (a_{i1}, \ldots, a_{in})^T$.

Consider a system of n linear equations in n variables:

$$\sum_{j=1}^{n} a_{ij} x_j \;=\; b_i \;\text{ for }\; 1 \le i \le n, \;\text{ or, in matrix form: }\; Ax \;=\; b. \tag{2.1}$$

We shall assume throughout that (2.1) is consistent and that A does not contain a row of zeros. For $p \ge 1$, we define a diagonal matrix $D = \mathrm{diag}(1/\|a_1\|_p, \ldots, 1/\|a_n\|_p)$. The geometrically-scaled system is defined as

$$DAx \;=\; Db. \tag{2.2}$$

In some algorithms, GS(p) is an inherent step in the following sense: either the scaling is executed at the beginning as an intrinsic part of the algorithm, or, executing the algorithm produces identical results to those obtained when GS(p) is done at the beginning. As an example, it is easy to see that GS(2) is inherent in the Kaczmarz algorithm (KACZ) [9]. KACZ can be described geometrically as follows: starting from an arbitrary point $x^0 \in \mathbb{R}^n$ as the initial iterate, KACZ projects the current iterate orthogonally towards a hyperplane defined by one of the equations. The hyperplanes are chosen in cyclic order.

The tests were run using the AZTEC software library [19], which includes a wide range of algorithms and preconditioners, suitable for sequential and parallel implementations. Geometric scaling with the L_1-norm is a built-in option in AZTEC (where it is called "row-sum scaling"). As mentioned we used Bi-CGSTAB and GMRES, with and without ILUT. Restarted GMRES was used with Krylov subspace size of 10. In AZTEC, GMRES is implemented with a double classical Gram-Schmidt orthogonalization step. ILUT was implemented with AZTEC's default parameters of drop tolerance = 0 and fill-in = 1. These parameters are not necessarily optimal for the tested examples, but since our purpose was to demonstrate the general usefulness of GS, we stuck with a commonly-used restart value for GMRES and AZTEC's default values for ILUT. No doubt, better convergence properties and runtimes would be obtained if these parameters were fine-tuned to each problem. The eigenvalue computations were done with the LAPACK linear algebra package [10].

We also experimented with the option of multiplying the system matrix A by $D^{\frac{1}{2}}$ on the left and on the right, resulting in the system $D^{\frac{1}{2}}AD^{\frac{1}{2}}y = D^{\frac{1}{2}}b$, with $x = D^{\frac{1}{2}}y$. In one problem, the results were somewhat poorer than those obtained with GS, but otherwise, the results were similar. Note that this option is computationally more expensive. Gambolati et al. [3] remark that when using ILU(0), the iteration matrix of Bi-CGSTAB on a diagonally scaled system (on the left and/or the right) is theoretically the same as the one with ILU(0) on the original matrix. In our experiments, the effect of ILUT was identical to that of ILU(0), so in theory, Bi-CGSTAB with ILUT on the original and the scaled systems should have behaved similarly. However, our experiments indicated that using GS produces much better numerical results, and this also holds for GMRES with ILUT. This is probably due to the fact that the scaled system does not have very large differences between coefficients, so it is much less prone to roundoff errors.

2.1 Setup of the numerical experiments

In two dimensions, the general form of the second-order differential equations in this study was

$$\frac{\partial}{\partial x}(a(x,y)u_x) + \frac{\partial}{\partial y}(b(x,y)u_y) + \cdots \;=\; F,$$

where a and b are given functions of x and y, "$\cdots$" stands for lower-order derivatives, and F is a prescribed RHS. In three dimensions, there are three given functions a, b, c of x, y, z, and a second order partial derivative w.r.t. z was also included. Boundary conditions were either Dirichlet or mixed Dirichlet and Neumann. The regions were taken as the unit square or the unit cube. The discretization of the second-order derivatives at a given grid point (i, j) was done

using central differences; e.g., $\frac{\partial}{\partial x}(au_x)$ was approximated as

$$\begin{aligned}\frac{\partial}{\partial x}(au_x)_{i,j} &= \left((au_x)_{i+\frac{1}{2},j} - (au_x)_{i-\frac{1}{2},j}\right)/\Delta x \\ &= \left(a_{i+\frac{1}{2},j}(u_{i+1,j}-u_{i,j})/\Delta x - a_{i-\frac{1}{2},j}(u_{i,j}-u_{i-1,j})/\Delta x\right)/\Delta x \\ &= \left(-(a_{i+\frac{1}{2},j}+a_{i-\frac{1}{2},j})u_{i,j} + a_{i+\frac{1}{2},j}u_{i+1,j} + a_{i-\frac{1}{2},j}u_{i-1,j}\right)/\Delta x^2\end{aligned}$$

All problems were discretized with equally-spaced grids, and the initial value was taken as $u^0 = 0$. The tests were run on a Pentium IV 2.8GHz processor with 3GB memory, running Linux.

2.2 Stopping tests

There are several stopping criteria which one may apply to iterative systems. Our stopping criterion was to use the relative residual: $\|b-Ax\|/\|b-Ax^0\| < \varepsilon$, where ε was taken as 10^{-4}, 10^{-7} and 10^{-10}. In some of the cases, this was not attainable. Since this stopping criterion depends on the scaling of the equations, we always made this test on the geometrically-scaled system using the L_2-norm. In order to limit the time taken by the methods implemented in AZTEC, the maximum number of iterations was set to 10,000. The AZTEC library has several other built-in stopping criteria: numerical breakdown, numerical loss of precision and numerical ill-conditioning. In the results, we denote the relative residual by rel-res, and non-convergence by "no conv."

One should note that the test for numerical breakdown in AZTEC uses the machine precision `DBL_EPSILON`, and this may result in a premature notice of numerical breakdown in some cases. To get around this problem, we multiplied the variable `brkdown_tol` in the Bi-CGSTAB algorithm by some small factor, such as 10^{-16}. (`brkdown_tol` is normally set equal to `DBL_EPSILON`, which is approximately 2.22×10^{-16} on our machine).

3 Problem 1

Problem 1 is taken from Saad [16, §3.7, problem F2DB]. It is a two-dimensional PDE

$$-\frac{\partial}{\partial x}(au_x) - \frac{\partial}{\partial y}(bu_y) + \frac{\partial(du)}{\partial x} + \frac{\partial(eu)}{\partial y} = h,$$

where

$$a(x,y) = b(x,y) = \begin{cases} 10^3 & \text{if } \frac{1}{4} < x,\ y < \frac{3}{4}, \\ 1 & \text{otherwise} \end{cases}$$

and $d(x,y) = 10(x+y)$ and $e(x,y) = 10(x-y)$. The Dirichlet boundary condition $u = 0$ is used on the boundary, and the RHS h is immaterial since the vector b of the linear system is chosen as $b = Ae$, where A is the system matrix and $e = (1,\ldots,1)^T$.

Table 3.1 shows the number of iterations and runtimes of the various algorithm/preconditioner methods, with and without GS, on a grid of 128×128. In cases of stagnation, the table shows the

relative residual that was achieved before stagnation. Note that GS enables the convergence of Bi-CGSTAB without ILUT; this is useful in a parallel setting, because Bi-CGSTAB is inherently parallel but ILUT is not an ideal parallel preconditioner. We can also see that GS is slightly helpful to Bi-CGSTAB with ILUT, and, for low-accuracy, also to GMRES with and without ILUT. Another observation is that when GMRES (with and without ILUT) stagnates before reaching the prescribed convergence goal, GS postpones the stage at which stagnation sets in, and enables convergence to a level that is acceptable for most practical applications.

	No. of iterations and time (in sec.)		
Method	rel-res $= 10^{-4}$	rel-res $= 10^{-7}$	rel-res $= 10^{-10}$
Bi-CGSTAB	no conv.	no conv.	no conv.
with GS	91 (0.30)	299 (0.99)	361 (1.19)
Bi-CGSTAB+ILUT	31 (0.23)	107 (0.67)	142 (0.88)
with GS	30 (0.23)	90 (0.59)	130 (0.81)
GMRES	converged to 3.8×10^{-2}		
with GS	265 (0.85)	converged to 1.1×10^{-5}	
GMRES+ILUT	converged to 3.9×10^{-3}		
with GS	39 (0.23)	converged to 1.1×10^{-5}	

Table 3.1: No. of iterations and runtimes for Problem 1. Grid size = 128×128.

Three additional experiments, based on Problem 1, were also done:

1. The values of $a(x,y)$ and $b(x,y)$ were increased to 10^4 in the inner square. The results were very similar to those shown in Table 3.1, but with slightly increased runtimes in the higher-accuracy cases.
2. A "continuous" case: the values of $a(x,y)$ and $b(x,y)$ were taken as 1 throughout the unit square. Here, all the methods converged, and GS made very little difference.
3. A second continuous case, with $a(x,y) = b(x,y) = 1000$. Table 3.2 shows the results of this experiment. As can be seen, GS made very little difference.

	No. of iterations and time (in sec.)		
Method	rel-res $= 10^{-4}$	rel-res $= 10^{-7}$	rel-res $= 10^{-10}$
Bi-CGSTAB	121 (0.41)	224 (0.76)	286 (0.96)
with GS	123 (0.42)	217 (0.74)	261 (0.88)
Bi-CGSTAB+ILUT	39 (0.28)	61 (0.41)	85 (0.56)
with GS	40 (0.29)	60 (0.40)	85 (0.56)
GMRES	1365 (4.31)	3704 (11.69)	6045 (19.19)
with GS	1356 (4.28)	3582 (11.45)	5722 (18.16)
GMRES+ILUT	140 (0.69)	359 (1.69)	579 (2.67)
with GS	140 (0.69)	359 (1.69)	579 (2.67)

Table 3.2: No. of iterations and runtimes for a continuous version of Problem 1, with $a = b =$ 1000 everywhere.

We turn now to studying the effect of GS on the distribution of the eigenvalues. For Problem 1, this was done with a grid size of 40×40 for the purpose of a clear presentation. Table 3.3

shows the values of the minimum and maximum eigenvalues and the condition number, for the original, the scaled and the continuous ($a = b = 1000$) cases of Problem 1. Also shown for each case is the number of eigenvalues in the first interval (out of 100) in a histogram of the eigenvalue distribution. We can see that while the condition number is not changed much by the scaling, the number of eigenvalues in the first interval is reduced very significantly.

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	No. of eigenvalues in first interval
Original	1.87E-2	7.96E+3	4.25E+5	1126
With GS	7.07E-6	1.78E+0	2.52E+5	4
Cont. coef.	1.17E+1	8.00E+3	6.81E+2	8

Table 3.3: Basic eigenvalue information for Problem 1.

Figure 3.1 shows the distribution of the eigenvalues of Problem 1 at the full scale (left) and with a zoom (right). We can see that there is a congestion of eigenvalues close to zero. Figure 3.2 shows the distribution of the eigenvalues of Problem 1 after geometric scaling with the L_1-norm (left) and the L_2 (right). Both scalings produce fairly similar distributions; however we will present results throughout the paper for the L_2 scaling. As to the continuous version of the problem, its eigenvalues were all real and distributed evenly in the range $(0, 8000)$.

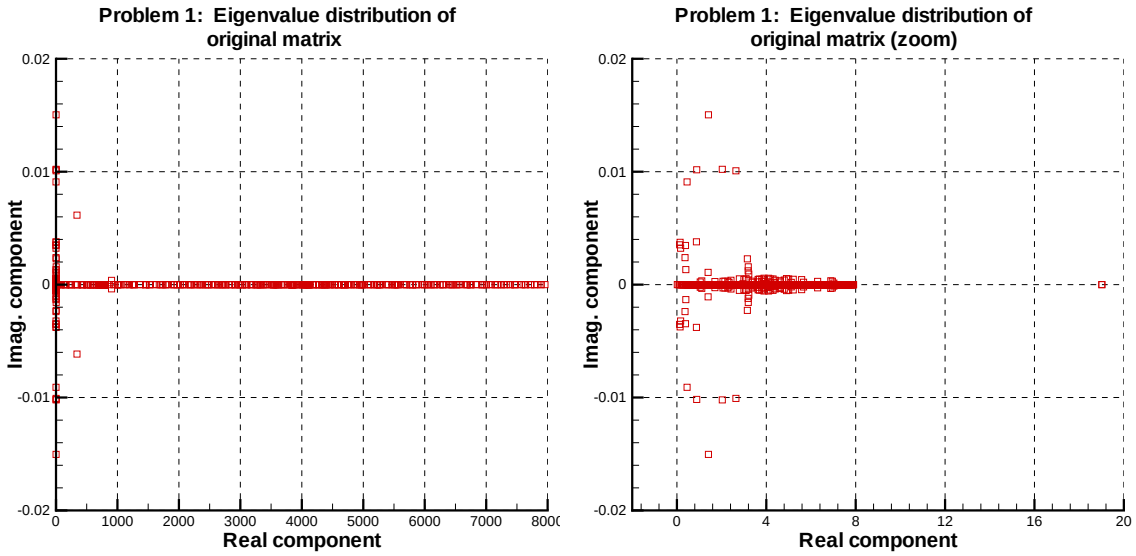

Figure 3.1: Eigenvalue distribution of the original matrix of Problem 1, with a zoom on the range 0–20.

Figure 3.3 shows a histogram of the eigenvalue distributions for the original Problem 1, for the geometrically scaled problem, and for the continuous case with $a = b = 1000$.

4 Problem 2

Problem 2 is based on Example 2 from van der Vorst [21], to which we added convection terms. This example demonstrates that GS also works with mixed Dirichlet and Neumann boundary

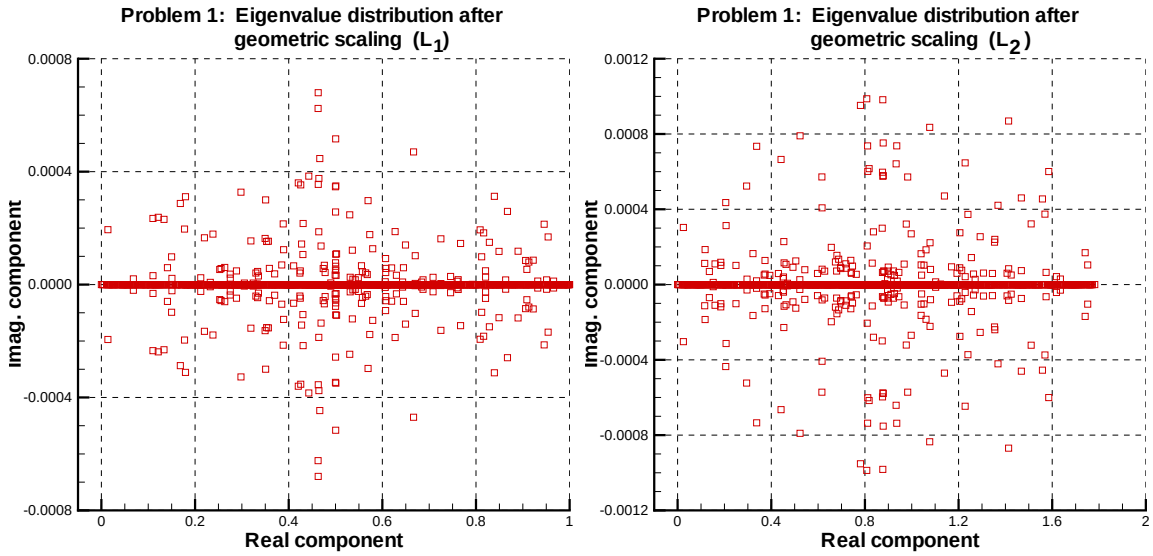

Figure 3.2: Eigenvalue distribution for Problem 1 after geometric scaling with the L_1- and the L_2-norms.

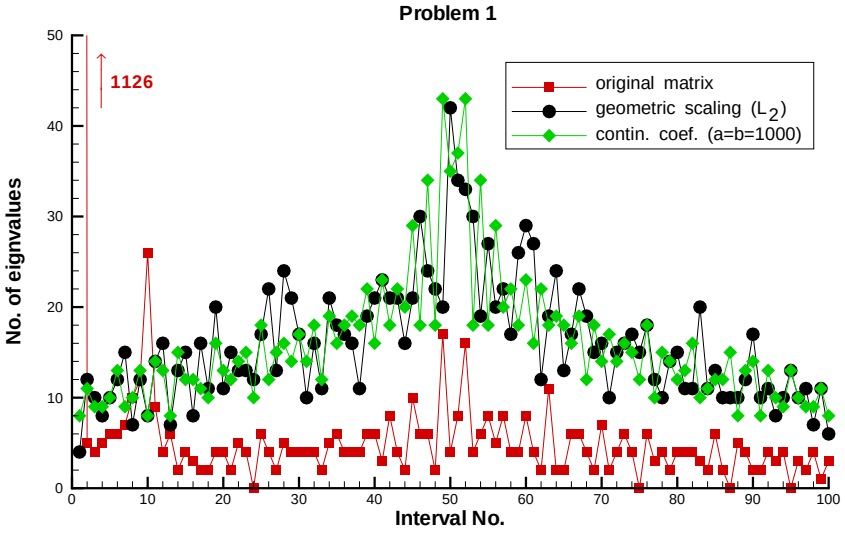

Figure 3.3: Histogram of eigenvalues for Problem 1, for the original matrix, the geometrically scaled matrix, and for a variation of Problem 1 with continuous coefficients ($a = b = 1000$).

conditions. The governing PDE is

$$-\frac{\partial(D(x,y)u_x)}{\partial x} - \frac{\partial(D(x,y)u_y)}{\partial y} + au_x + bu_y = 1,$$

with $D(x,y)$ and boundary conditions as shown in Figure 4.4. The convection terms were $a = b = 200$. In the original example, the internal value of D is 10^3, but we also tested internal values of 10^4 and 10^5. Also, the internal square in our case is somewhat smaller. The unit square was discretized with a grid size of 150×150. Together with the boundary equations, the system consisted of 22,952 equations. The resulting system is indefinite, with eigenvalues in the four quadrants of the imaginary plane.

Tables 4.4, 4.5 and 4.6 show the number of iterations and runtimes of the various algorithm

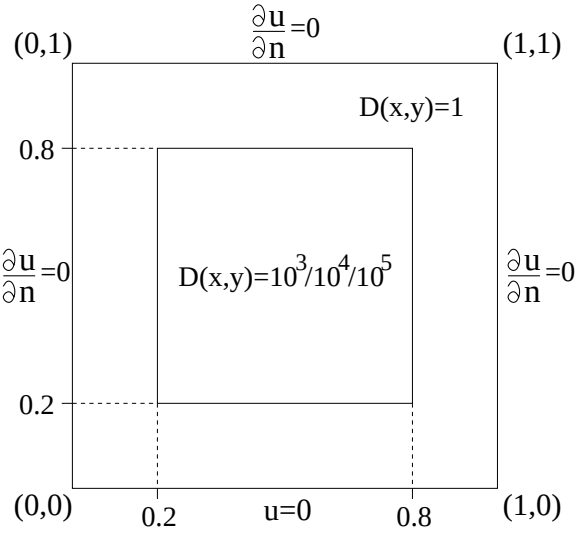

Figure 4.4: Description of Problem 2.

and preconditioner combinations, with and without GS, for $D = 10^3$, $D = 10^4$, and $D = 10^5$, respectively. When there was no convergence to the prescribed goal, the table shows the relative residual that was achieved.

	No. of iterations and time (in sec.)		
Method	rel-res $= 10^{-4}$	rel-res $= 10^{-7}$	rel-res $= 10^{-10}$
Bi-CGSTAB	no conv.	no conv.	no conv.
with GS	591 (2.13)	1025 (3.69)	conv. to 7.23×10^{-10}
Bi-CGSTAB+ILUT	92 (0.75)	139 (1.09)	215 (1.68)
with GS	88 (0.72)	125 (1.00)	324 (2.45)
GMRES	converged to 6.15		
with GS	converged to 2.145×10^{-3}		
GMRES+ILUT	converged to 0.927		
with GS	3361 (19.86)	converged to 3.67×10^{-5}	

Table 4.4: No. of iterations and runtimes for Problem 2 (internal $D = 10^3$).

	No. of iterations and time (in sec.)		
Method	rel-res $= 10^{-4}$	rel-res $= 10^{-7}$	rel-res $= 10^{-10}$
Bi-CGSTAB	no conv.	no conv.	no conv.
with GS	589 (2.01)	707 (2.41)	1493 (5.37)
Bi-CGSTAB+ILUT	96 (0.76)	192 (1.44)	255 (1.86)
with GS	79 (0.64)	121 (0.92)	163 (1.19)
GMRES	converged to 1.84		
with GS	converged to 2.19×10^{-4}		
GMRES+ILUT	converged to 0.924		
with GS	converged to 1.90×10^{-4}		

Table 4.5: No. of iterations and runtimes for Problem 2 (internal $D = 10^4$).

	No. of iterations and time (in sec.)		
Method	rel-res $= 10^{-4}$	rel-res $= 10^{-7}$	rel-res $= 10^{-10}$
Bi-CGSTAB	no conv.	no conv.	no conv.
with GS	224 (0.82)	730 (2.65)	conv. to 1.69×10^{-9}
Bi-CGSTAB+ILUT	96 (1.15)	192 (1.56)	conv. to 5.16×10^{-10}
with GS	21 (0.23)	125 (0.98)	217 (1.68)
GMRES	converged to 1.37		
with GS	357 (1.38)	converged to 2.19×10^{-5}	
GMRES+ILUT	converged to 0.90		
with GS	40 (0.32)	converged to 1.89×10^{-5}	

Table 4.6: No. of iterations and runtimes for Problem 2 (internal $D = 10^5$).

The tables show that GS was helpful in all cases, with the exception of Bi-CGSTAB with ILUT in Table 4.4, where it increased the number of iterations required in the highest accuracy case. A close examination of this case revealed that starting from iteration 220, BiCGSTAB with ILUT (after GS) was extremely oscillatory, with fluctuations of up to three orders of magnitude in the relative residual; the required convergence goal of 10^{-10} was almost attained at 225 iterations.

The three tables show that the behavior of GS is not necessarily "smooth": consider the results of Bi-CGSTAB with ILUT for rel-res = 10^{-10}: we can see that the goal was not reached for $D = 10^3$ and $D = 10^5$, but it was reached for $D = 10^4$. The explanation for this is the large number of iterations required by Bi-CGSTAB in this case, as seen in Table 4.5 and the oscillatory nature of Bi-CGSTAB. The accumulated roundoff error is sometimes too large to reach the convergence goal, and while the updated error measured by the algorithm showed convergence, the true relative residual did not always decrease sufficiently.

Another interesting point is the behavior of GS with GMRES and ILUT in the low-accuracy case: for $D = 10^3$, there was convergence after many iterations, for $D = 10^4$ there was no convergence, and then for $D = 10^5$ there was a very fast convergence. In order to examine this behavior, we also tested this particular case with $D = 10^6$, and we noticed that starting $D = 10^4$, GS improved the convergence of GMRES by one order of magnitude for every increase in D by one order of magnitude. Not doubt, this phenomenon should be studied further. In any case, for the larger values of D, GS enabled the convergence of GMRES, with and without ILUT, to practical levels of accuracy.

For the eigenvalue data, we discretized the domain by a grid of $40\times40\times40$. Table 4.7 provides the basic eigenvalue information for Problem 2, for the original ($D = 10^3$) and the scaled matrices. Also shown are the eigenvalues for a variation of Problem 2 with "continuous" coefficients, obtained by taking $D = 10^3$ throughout the unit square. The last column of the table shows the number of eigenvalues whose real part lies in the same interval as the origin, when the range of (the real parts) of the eigenvalues is divided into 100 intervals. We can see that GS reduces the condition number by three orders of magnitude, and the reduction is even greater w.r.t. the continuous case. Furthermore, the number of eigenvalues around the origin is reduced very significantly, even below that of the continuous case.

Figure 4.5 shows the distribution of the eigenvalues for Problem 2 with $D = 10^3$, and for the geometrically scaled problem. It can be seen that the eigenvalues of the original matrix are very concentrated around the origin. In the scaled matrix, there are very few eigenvalues around the

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	Eigenvalues around $x=0$
Original ($D=10^3$)	3.33E-1	1.34E+7	4.02E+7	1058
With GS	9.15E-5	1.78E0	1.95E+4	8
Continuous coef.	1.48E-2	1.34E+7	9.09E+8	130

Table 4.7: Basic eigenvalue information for Problem 2.

origin, but many eigenvalues have their real part around 0.6.

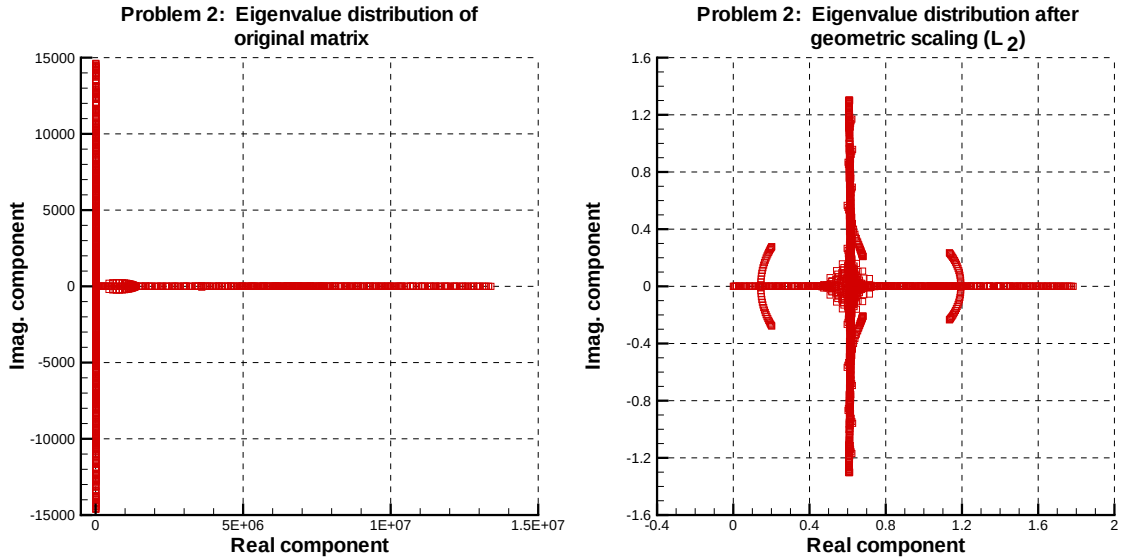

Figure 4.5: Eigenvalue distribution for Problem 2, for the original and the scaled cases.

Figure 4.6 provides another view of the eigenvalue distribution with a histogram of the real part of the eigenvalues for the original and the scaled matrices.

5 Problem 3

Problem 3 is also taken from van der Vorst [21, Example 4]. It describes a certain groundwater flow problem which leads to a nonsymmetric system, with a complex geometry and several jumps in the discontinuities of the equations. This problem is well-known for its difficulty. The basic equation is the following:

$$-\frac{\partial}{\partial x}(A(x,y)u_x) - \frac{\partial}{\partial y}(A(x,y)u_y) + B(x,y)u_x = F,$$

where $A(x,y)$ and F are taken as shown in Figure 5.7, and $B(x,y) = 2\exp(2(x^2+y^2))$. The Dirichlet boundary conditions are taken as shown in Figure 5.7. The unit square was discretized with a grid size of 128×128, resulting in 16,129 equations.

Nothing converged without ILUT, so we present in Table 5.8 only the results with ILUT, with and without GS. On this difficult problem, GS reduced the runtime of Bi-CGSTAB with ILUT

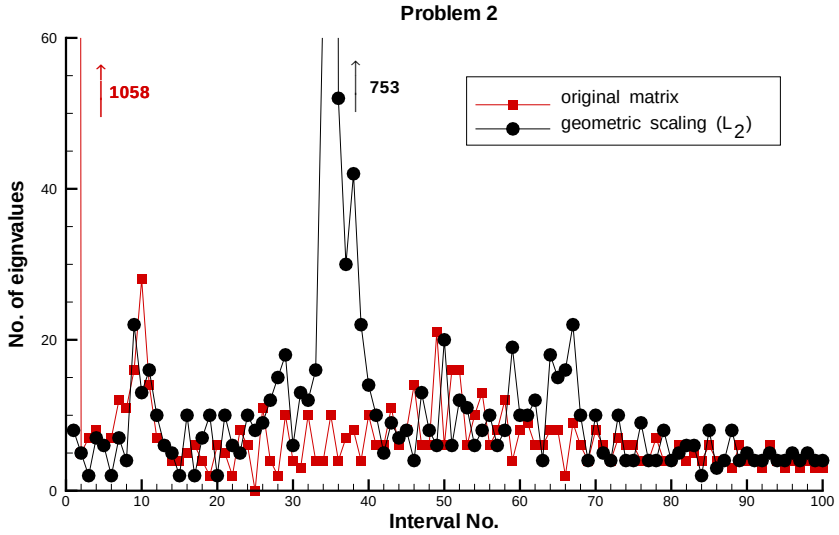

Figure 4.6: Histogram of the real part of the eigenvalues for Problem 2, for the original and the scaled cases.

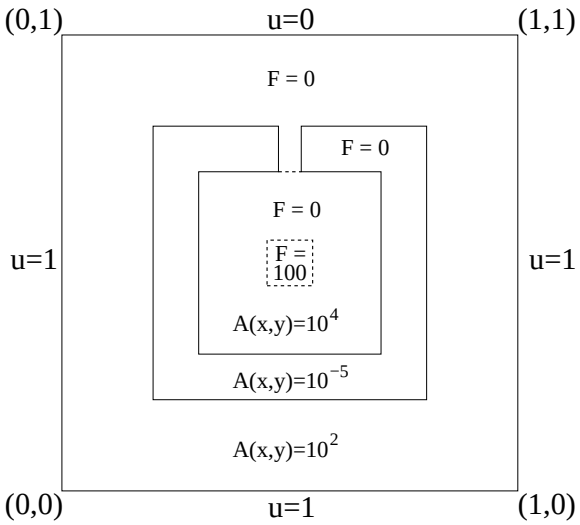

Figure 5.7: Description of Problem 3.

by about 25%. GS also enabled the convergence of GMRES with ILUT, though the achieved runtimes were an order of magnitude larger than those of the scaled Bi-CGSTAB with ILUT in the higher-accuracy cases.

	No. of iterations and time (in sec.)		
Method	rel-res $= 10^{-4}$	rel-res $= 10^{-7}$	rel-res $= 10^{-10}$
Bi-CGSTAB+ILUT	93 (0.60)	124 (0.79)	152 (0.95)
with GS	67 (0.44)	90 (0.59)	112 (0.72)
GMRES+ILUT	no conv.	no conv.	no conv.
with GS	338 (1.58)	1008 (4.61)	1683 (7.71)

Table 5.8: No. of iterations and runtimes for Problem 3.

For the eigenvalue data, we discretized Problem 3 with a grid of 40×40. The data for the original and the scaled matrices is summarized in Table 5.9, which also shows the number of eigenvalues in the first interval (out of 100) of the eigenvalue histogram. We can see that the condition number is decreased very significantly, and so is the number of eigenvalues in the first interval.

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	No. of eigenvalues in first interval
Original	6.97E-5	7.93E+4	1.14E+9	1284
With GS	1.21E-4	1.78E+0	1.47E+4	167

Table 5.9: Basic eigenvalue information for Problem 3.

Figure 5.8 shows the distribution of the eigenvalues for the original (with a zoom) and the scaled matrices of Problem 3.

Figure 5.9 is a histogram of the eigenvalues of the original and the scaled matrices. We can see from Figures 5.8 and 5.9 that in the original matrix, the eigenvalues are very concentrated around the origin, while in the scaled matrix, many are "pushed" away from the origin.

6 Problem 4

This problem is based on a three-dimensional problem from Graham and Hagger [8], to which we added convection terms. The differential equation is the following:

$$-\frac{\partial}{\partial x}(au_x) - \frac{\partial}{\partial y}(au_y) - \frac{\partial}{\partial z}(au_z) + du_x + eu_y + fu_z = 0,$$

where the domain is the unit cube, and $a(x,y,z)$ is defined as

$$a(x,y,z) = \begin{cases} D & \text{if } \frac{1}{3} < x,\ y,\ z < \frac{2}{3}, \\ 1 & \text{otherwise.} \end{cases}$$

The Dirichlet boundary conditions are prescribed with $u = 1$ on the $z = 0$ plane and $u = 0$ on the other boundaries. The convection terms were taken as $d = e = f = 100$, and two values of D were tested: $D = 10^4$ (as in the original problem) and $D = 10^6$. Two grids were tested: $40\times 40\times 40$ and $80\times 80\times 80$. The resulting linear systems are indefinite, with eigenvalues in the four quadrants of the imaginary plane. This problem will also be used to test the limit of usefulness of GS as the convection increases.

Due to the many different cases, the data for this problem is presented differently. Table 6.10 shows the number of iterations required for convergence, for a grid of $80\times 80\times 80$, with $D = 10^4$ and $D = 10^6$. We can see that GS enabled the convergence of Bi-CGSTAB without ILUT, and it speeded up Bi-CGSTAB with ILUT quite significantly. GS also enabled the convergence of GMRES, with and withou ILUT, for rel-res = 10^{-4}, and also for rel-res = 10^{-7} with $D = 10^6$.

Figures 6.10 and 6.11 show bar-plots of the runtimes for the two grid sizes, for $D = 10^4$ and $D = 10^6$, for the three levels of accuracy. In cases of stagnation, the figures show (in parentheses) the relative residual achieved before stagnation.

The results can be summarized as follows.

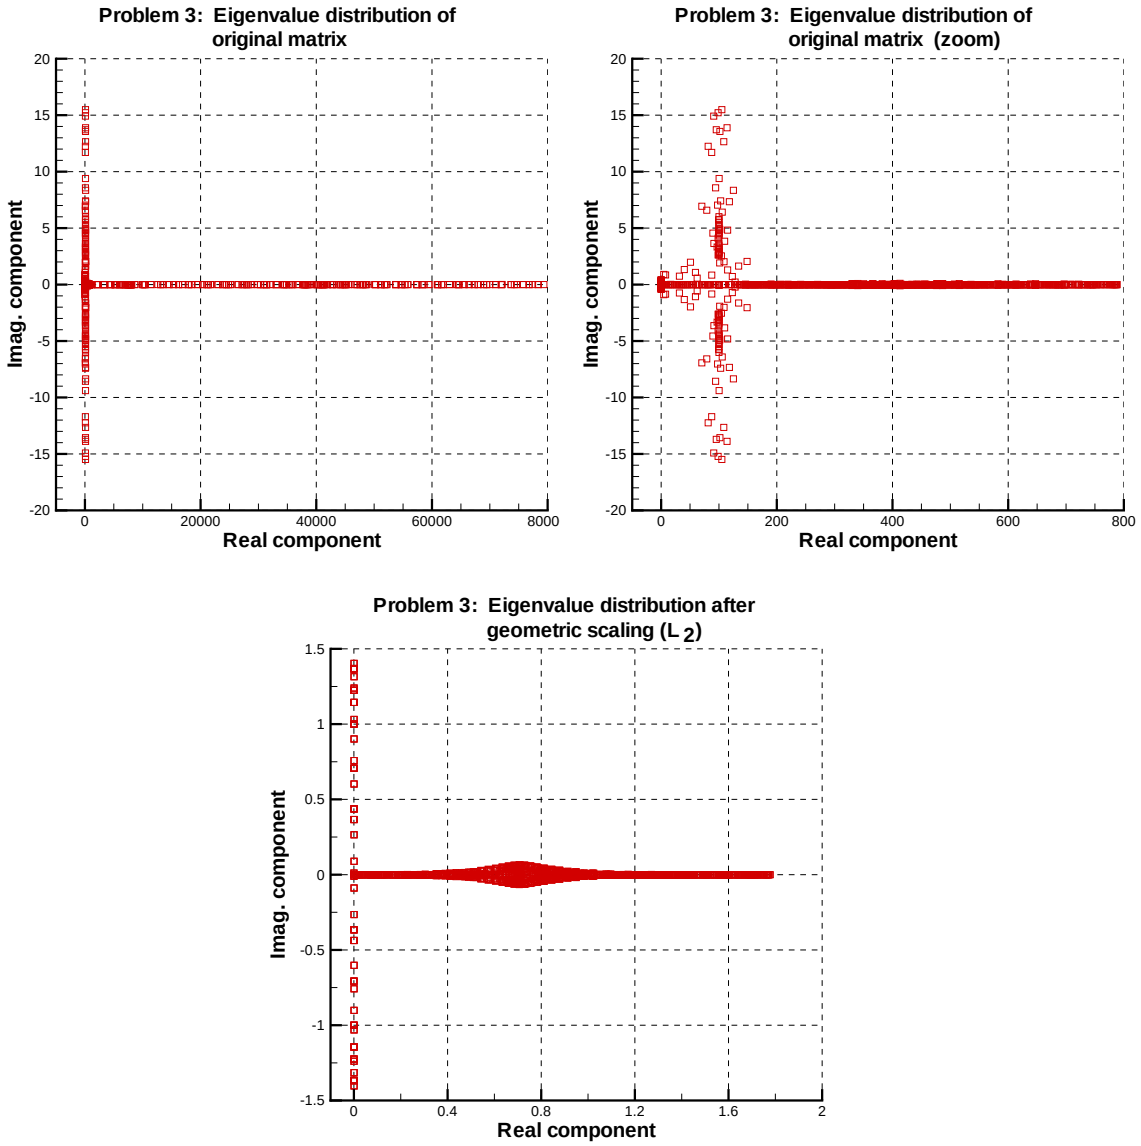

Figure 5.8: Eigenvalue distribution for Problem 3, with a zoom to the region 0–800, and the distribution for the scaled matrix.

- Bi-CGSTAB (without ILUT) and GMRES (with and without ILUT) did not converge in any of the cases.
- GS enabled the convergence of Bi-CGSTAB in all cases, and the convergence of GMRES (with and without ILUT) in the low-accuracy cases and in one mid-accuracy case.
- GS significantly improved the accuracy of GMRES (with and without ILUT) by postponing the stage at which stagnation set in, enabling convergence to much higher levels of accuracy.
- For the $80\times80\times80$-grid, GS speeded up the convergence of Bi-CGSTAB with ILUT quite significantly. This was also true for the coarser grid, but only in the low-accuracy case.
- In the above cases, GS (by itself) was generally a better preconditioner than ILUT for Bi-CGSTAB.
- In all cases with GS, the results for $D = 10^6$ were either similar to or better than those for $D = 10^4$.
- In summary, GS was beneficial in all cases to all the algorithm/ILUT combinations, and in

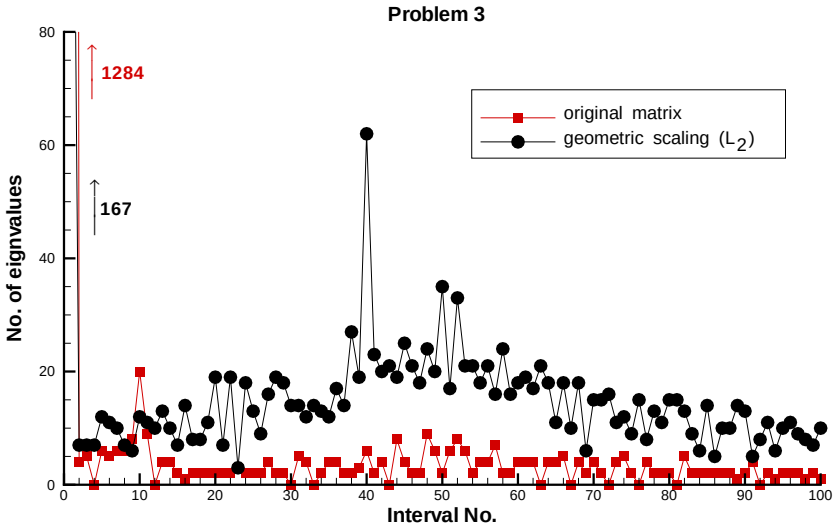

Figure 5.9: Histogram of the eigenvalues for Problem 3, for the original and the scaled cases.

	rel-res $= 10^{-4}$		rel-res $= 10^{-7}$		rel-res $= 10^{-10}$	
Method	$D = 10^4$	$D = 10^6$	$D = 10^4$	$D = 10^6$	$D = 10^4$	$D = 10^6$
Bi-CGSTAB	—	—	—	—	—	—
with GS	112	111	262	160	406	374
Bi-CGSTAB+ILUT	77	52	112	139	123	169
with GS	13	13	52	19	84	86
GMRES	—	—	—	—	—	—
with GS	208	207	—	286	—	—
GMRES+ILUT	—	—	—	—	—	—
with GS	28	28	—	46	—	—
Note: "—" denotes no convergence to the prescribed accuracy.						

Table 6.10: No. of iterations for Problem 4, with $D = 10^4$, $D = 10^6$, and grid $80 \times 80 \times 80$.

many cases, GS provided a very significant advantage.

Table 6.11 provides the basic eigenvalue information for Problem 4 with $D = 10^4$, for the original and the scaled matrices, and also for a continuous version with $D = 10^4$ everywhere. The grid size for this data was $12 \times 12 \times 12$. The the last column shows the number of eigenvalues whose real part lies in the same interval as the origin (out of 100 intervals). We can see that although GS doubled the condition number, it reduced the number of eigenvalues around the origin very significantly.

Figure 6.12 shows the distribution of the eigenvalues for the original Problem 4, and a zoom to the region $[-300, 300]^2$. We can see that the eigenvalues are very concentrated around the origin. Figure 6.13 shows the distribution of the eigenvalues for the scaled and the continuous ($D = 10^4$) cases. We can see that these distributions look quite similar, with eigenvalues concentrated around the perimeter.

Figure 6.14 provides a histogram of the real part of the eigenvalues of Problem 4, for the origi-

Results for $D = 10^4$ Results for $D = 10^6$

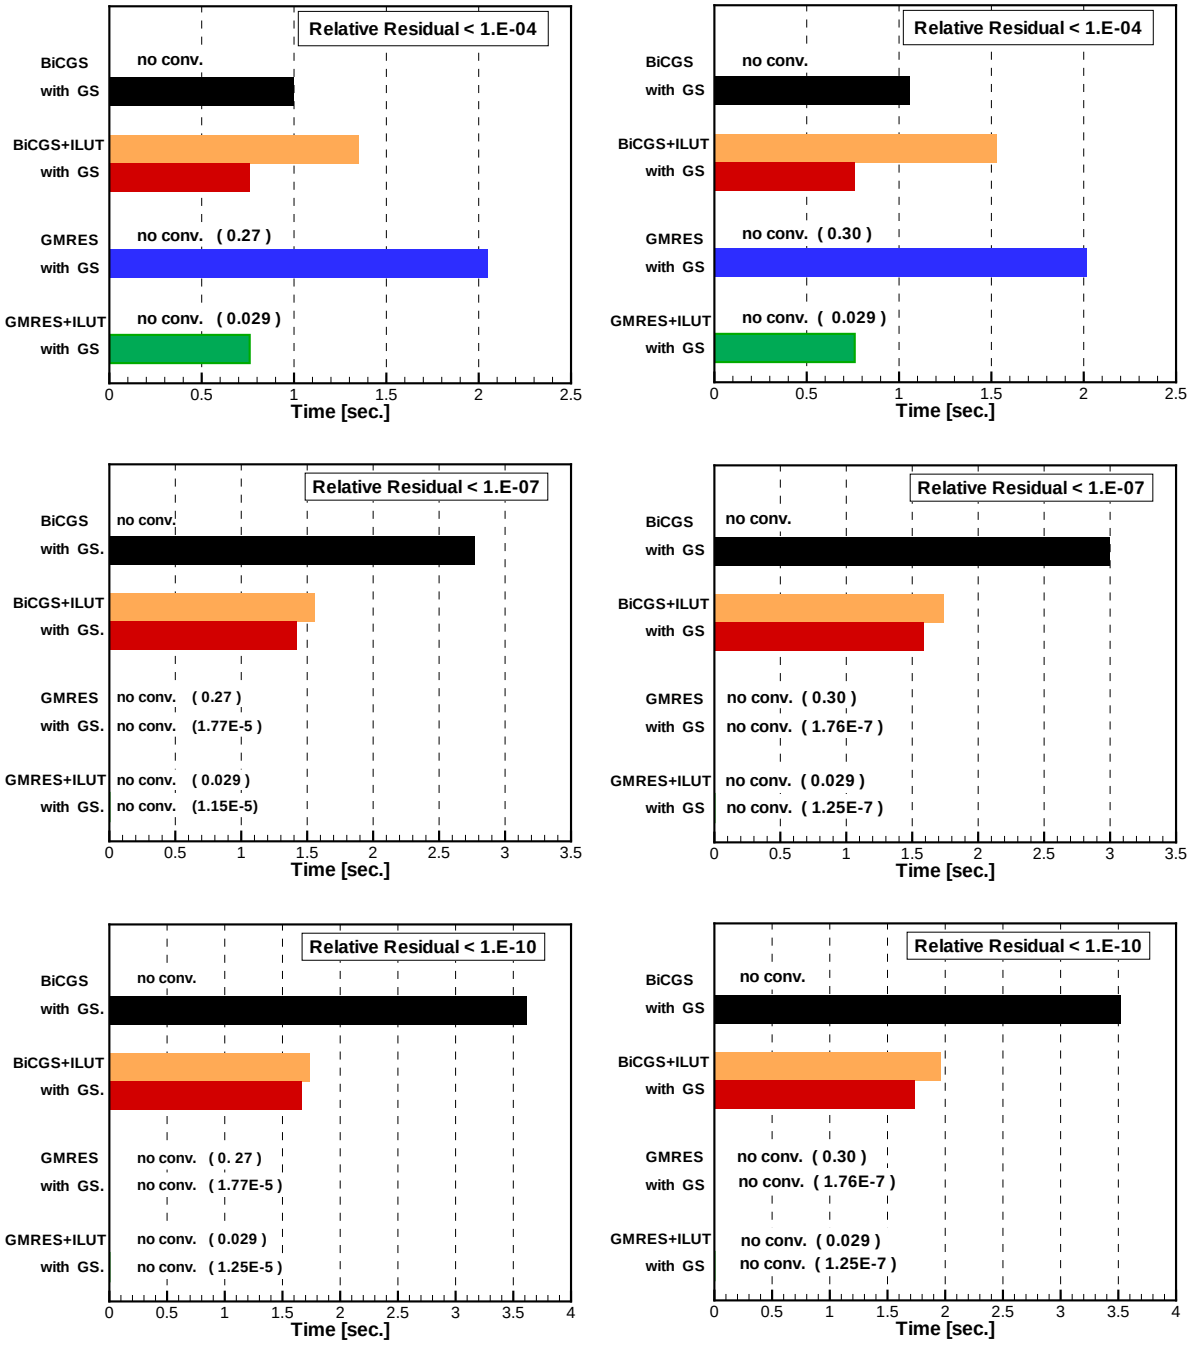

Figure 6.10: Runtimes and convergence status for Problem 4, grid = $40\times40\times40$.

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	No. of eigenvalues around $x=0$
Original	1.45E-2	2.93E+3	2.02E+5	1131
With GS	2.41E-6	1.00E+0	4.16E+5	25
Cont. coef. ($D=10^4$)	3.61E+0	6.18E+4	1.71E+4	42

Table 6.11: Basic eigenvalue information for Problem 4.

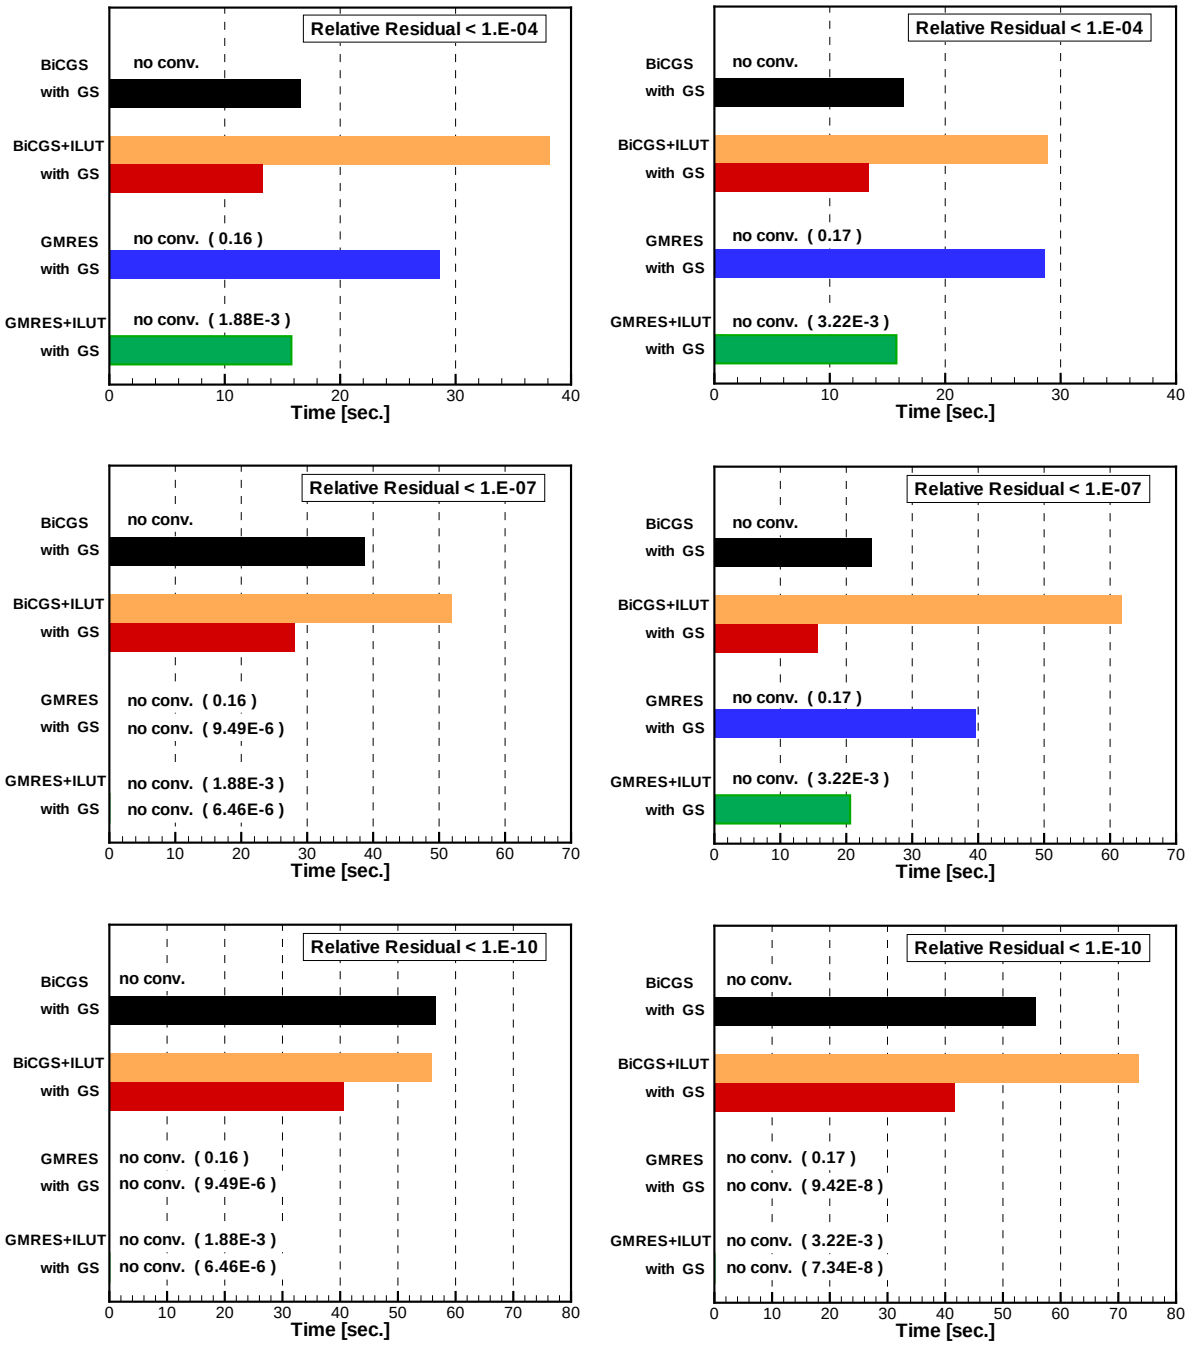

Figure 6.11: Runtimes and convergence status for Problem 4, grid = $80\times80\times80$.

nal, the scaled, and the continuous ($D = 10^4$) cases. The histogram provides an additional view of how the eigenvalues are distributed in the three cases.

GS is proposed as being useful for discontinuous coefficients when the convection terms are small to moderate. In order to demonstrate this, we considered Problem 4 with $D = 10^4$, a grid size of $40\times40\times40$, and varying convection terms. The previous results were based on convection terms of 100, so we also tested this case with convection terms of 200, 500, and 1000. The results, which are summarized in Table 6.12 below, show how the usefulness of GS

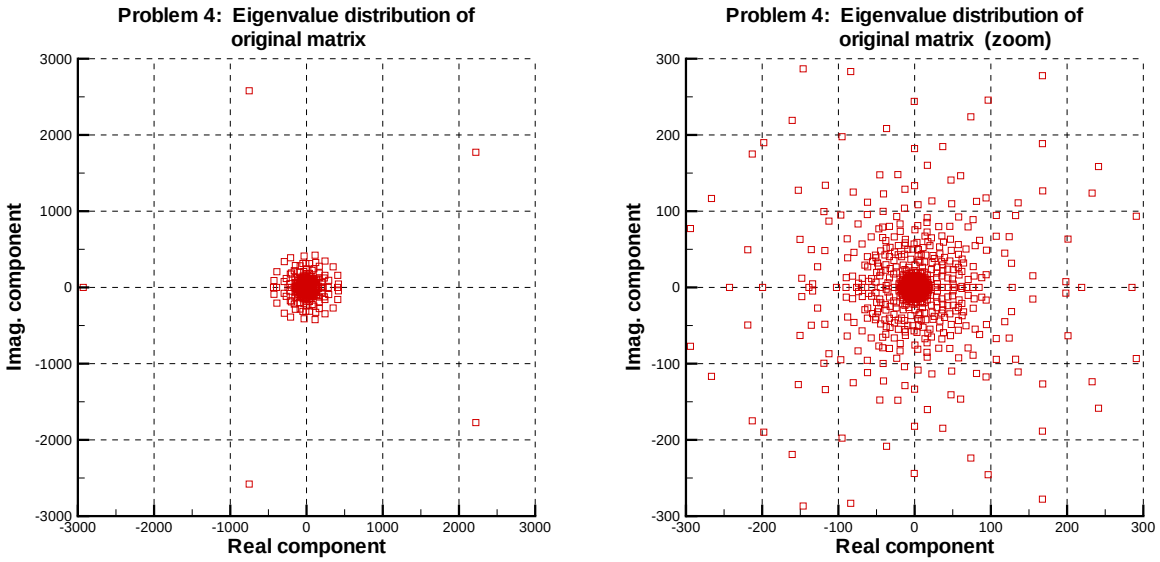

Figure 6.12: Eigenvalue distribution for Problem 4, with a zoom to the region $[-300, 300]^2$.

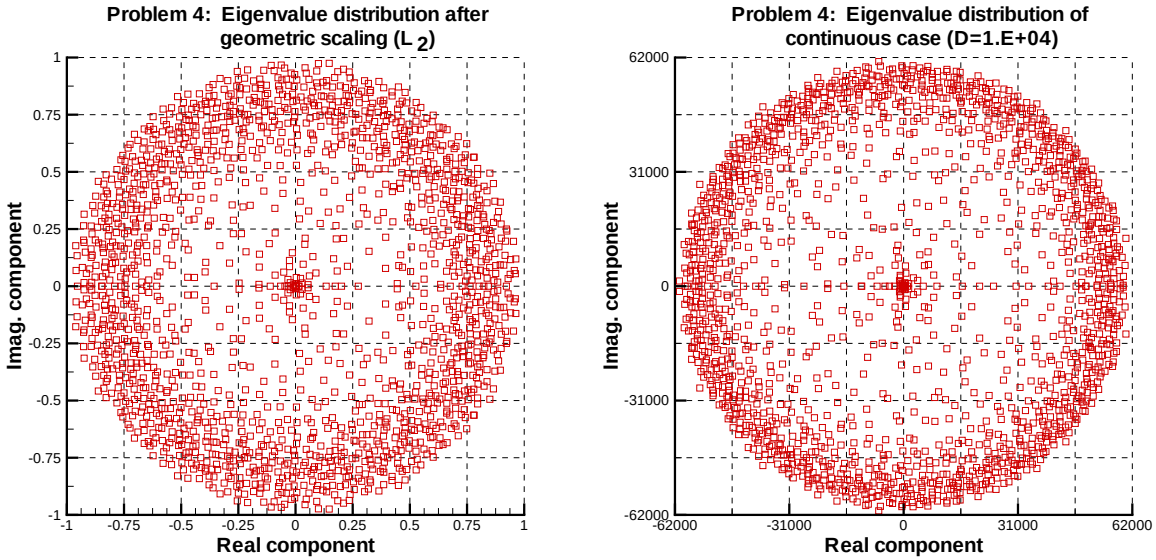

Figure 6.13: Eigenvalue distribution for Problem 4, for the scaled and the continuous cases.

degrades as the convection terms are increased.

Note in particular the difference between Bi-CGSTAB without ILUT and Bi-CGSTAB with ILUT: it turns out that without ILUT, GS enables convergence to 10^{-10} even with convection = 500, but with ILUT, degradation already appears with convection = 200; this degradation is independent of the use of GS.

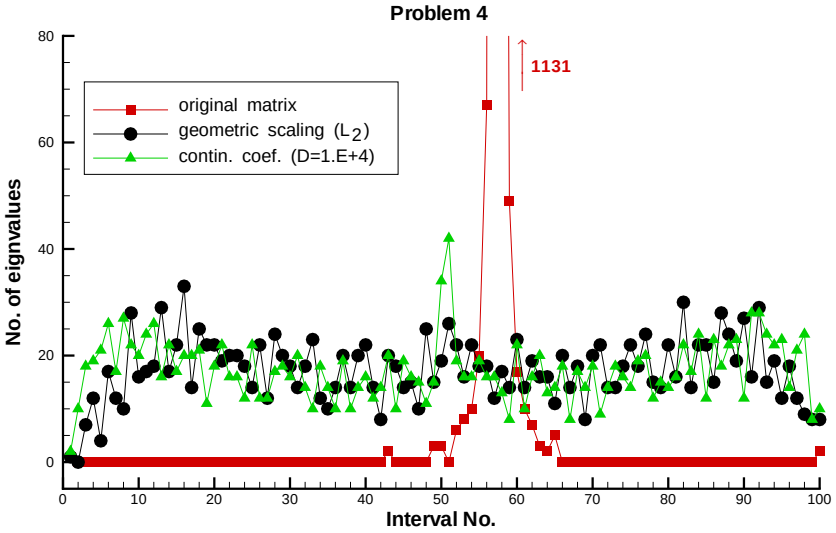

Figure 6.14: Histogram of the real part of the eigenvalues for Problem 4, for the original, the scaled, and the continuous cases.

Method / Convection:	**100**	**200**	**500**	**1000**
Bi-CGSTAB	—	—	—	—
with GS	10^{-10}	10^{-10}	10^{-10}	—
Bi-CGSTAB+ILUT	10^{-10}	10^{-4}	—	—
with GS	10^{-10}	10^{-4}	—	—
GMRES	—	—	—	—
with GS	10^{-4}	10^{-4}	—	—
GMRES+ILUT	—	—	—	—
with GS	10^{-4}	—	—	—
Note: '—' means no convergence. The numbers indicate which relative error goal was obtained.				

Table 6.12: Degradation of the various methods as the convection terms are increased.

7 Conclusions and further research

In this paper we have reported on a very simple technique for improving the convergence properties of algorithms on certain linear systems with discontinuous coefficients. The technique, called geometric scaling (GS), consists of simply dividing each equation by the L_p-norm of its vector of coefficients. GS, with $p = 2$, was tested on four nonsymmetric problems derived from convection-diffusion elliptic PDEs, with small to moderate convection terms. Bi-CGSTAB and restarted GMRES, with and without ILUT, were tested on the four problems with and without GS. Normally, such problems are solved by domain decomposition (DD) techniques, but these can be difficult to implement if the boundaries between the subdomains have a complicated geometry or the grid is unstructured. The problems were taken from (or based on) well-known examples from the literature, and included two- and three-dimensional cases, with Dirichlet and mixed Dirichlet/Neumann boundary conditions. One problem had a complicated geometry. Three different convergence goals were prescribed: relative residual $\leq 10^{-4}, 10^{-7}, 10^{-10}$.

Table 7.13 summarizes the convergence behavior of the four different algorithm/preconditioner combinations, with and without GS, on the four different problems, for the three convergence goals. The results indicate that GS was very useful in improving the convergence status of the different solution methods on the tested problems. With one exception, when the solution method converged on a problem with discontinuous coefficients, GS speeded up the convergence. Eigenvalue distribution maps show that when there is a strong concentration of values around the origin, GS reduces this concentration very significantly and "pushes" the values away from the origin.

Method	**Problem 1**	**Problem 2**	**Problem 3**	**Problem 4**
Bi-CGSTAB	$- \; - \; -$	$- \; - \; -$	$- \; - \; -$	$- \; - \; -$
with GS	$+ \; + \; +$	7.2×10^{-10}	$- \; - \; -$	$+ \; + \; +$
Bi-CGSTAB+ILUT	$+ \; + \; +$	$+ \; + \; *$	$+ \; + \; +$	$+ \; + \; +$
with GS	$+ \; * \; *$	$* \; * \; +$	$* \; * \; *$	$* \; * \; *$
GMRES	3.8×10^{-2}	1.37	$- \; - \; -$	0.27
with GS	1.1×10^{-5}	2.2×10^{-5}	$- \; - \; -$	1.8×10^{-5}
GMRES+ILUT	3.9×10^{-3}	0.90	$- \; - \; -$	0.29
with GS	1.1×10^{-5}	1.9×10^{-5}	$+ \; + \; +$	1.25×10^{-5}
Note: '$-$' means no convergence, '$+$' means convergence, '$*$' means better convergence. The numbers indicate the best relative error obtained.				

Table 7.13: Summary of convergence of the different methods on the four problems, for the three convergence goals ($10^{-4}, 10^{-7}$, and 10^{-10}).

This study is only a first report on this subject, and much work remains to be done. On the theoretical side, we have seen that GS improves the eigenvalue distribution, but further analysis is still needed. On the practical side, GS needs to be tested in conjunction with other algorithms and preconditioners, on various other difficult problems.

A natural question that arises is how does GS compare with DD techniques? Over the years, these methods have achieved a high level of efficiency, and a head-on runtime comparison is very much problem-dependent and a topic for further research. However, an even more interesting topic is the combination of the two methods: apply GS to the equations and then apply some standard DD method; hopefully, on some problems, GS could also improve the convergence properties of some DD approaches.

GS can also help the parallelization of solution methods in several ways. Firstly, GS itself is a parallel computational step. Secondly, it enables the partitioning of a domain into subdomains along boundaries that do not necessarily follow the physical boundaries of the problem; this way, better load balancing can be achieved. Thirdly, as seen in some of the cases, GS enabled the convergence of algorithms that are inherently parallel without the need for a preconditioner (such as ILUT) that may not be ideal for parallelism.

Another topic for further research is the application of the (sequential) CGMN algorithm [1, 6] and its block-parallel version CARP-CG [5] on problems that have discontinuous coefficients and are also strongly convection-dominated. These algorithms have already been shown to be very effective on convection-dominated problems and initial experiments with discontinuous coefficients have yielded good results. The reason for this may be that these two methods are essentially conjugate-gradient acceleration of the Kaczmarz algorithm, and as such, GS (with the L_2-norm) is already inherent in them.

References

[1] Å. Björck and T. Elfving. Accelerated projection methods for computing pseudoinverse solutions of systems of linear equations. *BIT*, 19:145–163, 1979.

[2] I. S. Duff and H. A. van der Vorst. Preconditioning and parallel preconditioning. Technical Report TR/PA/98/23, CERFACS, Toulouse, France, 1998.

[3] G. Gambolati, G. Pini, and M. Ferronato. Scaling improves stability of preconditioned CG-like solvers for FE consolidation equations. *International J. for Numerical & Analytical Methods in Geomechanics*, 27:1043–1056, 2003.

[4] R. Glowinski and M. F. Wheeler. Domain decomposition and mixed finite element methods for elliptic problems. In R. Glowinski, G. H. Golub, G. A. Meurant, and J. Périaux, editors, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, pages 144–172. SIAM, Philadelphia, PA, USA, 1988.

[5] D. Gordon and R. Gordon. CARP-CG: a robust and efficient parallel solver for linear systems, applied to strongly convection-dominated elliptic partial differential equations. Technical report, Dept. of Computer Science, University of Haifa, Haifa, Israel, Dec. 2008.

[6] D. Gordon and R. Gordon. CGMN revisited: robust and efficient solution of stiff linear systems derived from elliptic partial differential equations. *ACM Trans. on Mathematical Software*, 35(3):18:1–18:27, Oct. 2008.

[7] D. Gordon and R. Mansour. A geometric approach to quadratic optimization: an improved method for solving strongly underdetermined systems in CT. *Inverse problems in Science & Engineering*, 15(8):811–826, Dec. 2007.

[8] I. G. Graham and M. J. Hagger. Unstructured additive Schwarz-conjugate gradient method for elliptic problems with highly discontinuous coefficients. *SIAM J. on Scientific Computing*, 20(6):2041–2066, 1999.

[9] S. Kaczmarz. Angenäherte Auflösung von Systemen linearer Gleichungen. *Bulletin de l'Académie Polonaise des Sciences et Lettres*, A35:355–357, 1937.

[10] LAPACK – Linear Algebra Package. http://www.netlib.org/lapack.

[11] G. Meurant. *Computer Solution of Large Linear Systems*. Elsvier, Amsterdam, 1999.

[12] D. W. Peaceman and H. H. Rachford, Jr. The numerical solution of parabolic and elliptic differential equations. *Journal of the Society of Industrial and Applied Mathematics*, 3:28–41, 1955.

[13] A. Quarteroni and A. Valli. *Domain Decomposition Methods for Partial Differential Equations*. Oxford University Press, Oxford, 1999.

[14] J. R. Rice, P. Tsompanopoulou, and E. Vavalis. Interface relaxation methods for elliptic differential equations. *Applied Numerical Mathematics*, 32:219–245, 2000.

[15] Y. Saad. ILUT: a dual threshold incomplete ILU factorization. *Numerical Linear Algebra with Applications*, 1:387–402, 1994.

[16] Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, PA, 2nd. edition, 2003.

[17] Y. Saad and M. H. Schultz. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. on Scientific & Statistical Computing*, 7:856–869, 1986.

[18] B. Smith, P. Bjørstad, and W. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Diffferential Equations*. Cambridge University Press, Cambridge, UK, 1996.

[19] R. S. Tuminaro, M. A. Heroux, S. A. Hutchinson, and J. N. Shadid. AZTEC user's guide. Technical Report SAND99-8801J, Sandia National Laboratories, Albuquerque, New Mexico, 1999.

[20] A. van der Sluis. Condition numbers and equilibration of matrices. *Numerische Mathematik*, 14:14–23, 1969.

[21] H. A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. on Scientific & Statistical Computing*, 13:631–644, 1992.

[22] O. B. Widlund. On the effects of scaling of the Peaceman-Rachford method. *Mathematics of Computation*, 25(113):33–41, Jan. 1971.